

Partitioning of a Signal Detection Algorithm to a Heterogeneous Multicomputing Platform

Mr. Michael Vinskus, Principal Software Engineer

Mercury Computer Systems, Inc.

Phone: (978) 967-1653

FAX: (978) 224-0520

mvinskus@mc.com


Abstract:

This paper explores the tradeoffs involved in mapping a signals intelligence algorithm to general-purpose processor and field programmable gate array (FPGA) based technology. Specifically, a prototypical signal detection algorithm is described. This algorithm consists of an Fourier transform based frequency channelizer followed by a statistical signal detector.

The system examined consists of a single 14-bit real-valued input stream sampled at 100 MSa/s. With a Fourier transform overlap factor of 75%, this results in a total sustained bandwidth of 1 GB/s. The bandwidth is too large for a single commercial off-the-shelf (COTS) processor to get on and off board, leading to systems solutions using multiple processors.

The multiple-processor partitioning problem is looked at from both the time and frequency domains. For the example application, the strengths and weaknesses of each strategy are examined. The influence of the choice of processing platform on the partitioning affects the final solution as well. General-purpose processors typically run at very high speeds, but perform only a small number of operations per clock cycle. FPGAs, on the other hand, can perform thousands of operations per clock cycle, but operate with a slower clock frequency. These differences, as well as other system features such as the interprocessor communication subsystem, dramatically affect the viability of potential partitioning solutions.

It is shown that successful multiprocessor partitioning depends on the entire system. Of critical importance are the features and performance of the processing nodes and the interprocessor communications system. When the requirements are greater than a single aspect of the system can handle, this paper explores the possibility of utilizing excess capacity in other areas of the



system to balance the system loading. Finally, some of the issues that arise from extending the system to multiple antenna streams are also explored.

Figure 1: An example heterogeneous multicomputing computing platform.

Report Documentation Page			Form Approved OMB No. 0704-0188	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE 20 AUG 2004	2. REPORT TYPE N/A	3. DATES COVERED -		
4. TITLE AND SUBTITLE Partitioning of a Signal Detection Algorithm to a Heterogeneous Multicomputing Platform			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Mercury Computer Systems, Inc.			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited				
13. SUPPLEMENTARY NOTES See also ADM001694, HPEC-6-Vol 1 ESC-TR-2003-081; High Performance Embedded Computing (HPEC) Workshop (7th)., The original document contains color images.				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON




Figure 2: Some of the issues associated with heterogeneous multicompacting applications.



Partitioning of a Signal Detection Algorithm to a Heterogeneous Multicomputing Platform


Michael Vinskus
Mercury Computer Systems, Inc.

High Performance Embedded Computing (HPEC) Conference
September 23, 2003

The Ultimate Performance Machine

System Description


- The goal of the acquisition system is to detect the presence of new signals in the environment in a timely manner so that they may be identified and exploited
 - A homogeneous system based on general purpose processors can be used to solve the problem, but advances in FPGAs allow for a higher processing density, especially in the channelizer
 - Current FPGAs can offer over a 10x computational density improvement over general purpose processors for certain applications. Unfortunately, most communication fabrics have not scaled at the same rate



**How can the I/O issue be managed?
By partitioning across multiple processing elements!**


- Typical acquisition systems place the delay memory after the analog to digital converter. For demodulation, a digital down converter(DDC) is used to heterodyne and filter the delayed data stream
- When the number of signals to demodulate becomes very large (>100), the typical DDC-based approach becomes cumbersome. FFT processing can allow the simultaneous down conversion and filtering of thousands of signals. The downside is an increase in the amount of memory needed to give the same time delay

The Channelizer




- The channelizer decomposes the input sample stream into frequency channels by performing overlapped and windowed, short time Fourier transforms on the input data stream



- Parameters:
 - 16384 point real DFT
 - 4:1 overlap and windowing ($P = 4$)
 - 16384 input sample maximum latency requirement
 - Output 8192 bins, complex, 24 bits per component (1200 MB/s)



- The channelizer throughput is greater than most interconnect fabrics can support
- Need to partition the problem



Detection Processing




- The new energy alarm detects the presence of “new” signals and performs some rudimentary external measurements
- High-input bandwidth and computational requirements necessitate partitioning
- Partitioning strategies
 - Commutation
 - Time domain broadcast
 - Frequency division

Commutator Partitioning

- Each processing partition generates the entire frequency sweep for a range of time slices

- In this case, the channelizer is split into four partitions
- Each partition processes a contiguous segment of the input data stream
- By partitioning the problem in this manner, extra overhead to handle the segment boundaries is incurred
- Also, the system latency increases due to time expansion nature of the commutation process


- In this example, $P-1$ old input data blocks need to be received and $P-1$ blocks need to be sent
- Total I/O overhead is $2(P-1)$ blocks for transfers greater than $2P-1$ blocks of data



Commutator Implementation


- To meet the one block latency requirement, only 1/4 of a block of new data is passed to each partition; 3/4 of the data comes from the other partitions. This results in the same data block being transferred an extra 3 times!

- For this example, the input bandwidth increases from 200 MB/s to 800 MB/s. We're going the wrong way!



Time Domain Broadcast


- Each partition receives the entire data stream
- No extra I/O overhead is incurred
- Partitioning does not add latency




- No communication between channelizer partitions
- Still need to send all of the channelizer output data to the detector processors
- If there are a sufficient number of detector processing elements allocated and they are located correctly in the fabric, then I/O bottlenecks can be avoided

Broadcast Implementation

- Each fabric connection is a 266 MB/s half duplex link to the crossbar. Typical performance is around 250 MB/s
- To accommodate the channelizer output rate, the output stream must be split across multiple connections
- This complicates the data flow and fully utilizes the fabric I/O capacity in many places




Frequency Division Partitioning





- Able to pipeline the channelizer and detection processing with minimal inter-processor communication

- Each node generates a subset of the frequency bins for all time slices
- Each partition needs all of the time series data for each sweep



Frequency Division Mapping


- Each partition performs the front part of the FFT computations. This is inefficient, but the I/O issues are simpler
- Uniform communication between partitions at partition boundaries only
 - Single element communication between channelizer partitions
 - Similar communication between detector partitions!
- **So why the unusual allocation of the frequency bins?**



Frequency Division Rationale

- The algorithm exploits the efficient computation of the DFT of a **2N-point, real-valued, input sequence with an N-point complex transform**
- Since the input data is real, **only one half of the output spectrum needs to be computed**

$$G(k) = \frac{1}{2} \left\{ X(k) + X^*(N-k) - j \cdot W_{2N}^k \cdot (X(k) - X^*(N-k)) \right\} \quad \text{for } k = 0, 1, \dots, N-1.$$




By processing a range of $X(k)$, $X(N-k)$ frequency pairs, twice as many output points can be calculated with only two extra additions!


Frequency Division Approach

- Redundant computations are performed in each channelizer partition to reduce the system I/O requirements
- Unlike the time partitioning case, most of the I/O movement occurs locally in the fabric

Exchanges processing capacity for reduced I/O complexity



Multi-Antenna Implications



- Adds an extra data dimension to the problem
- Typical antenna arrays consist of 4 - 8 antenna elements

- The detection algorithms are different than the single antenna case. Typically, eigenspace methods are employed
 - The channelizer is similar between the single and multi-antenna cases
 - Detection processing requires the time series data for each frequency bin
 - Log magnitude computation is not required
 - Eigenvalues and eigenvectors are computed for each bin, across all antennas
- As shown in the previous example, frequency domain partitioning has advantages when I/O bound conditions occur
 - The majority of the data movement is local
 - High-speed local interconnects may be used instead of the fabric connections